ENQPY

Toy Run — Empirical Demonstration of the Ideal Configuration

R = 2 (binary) · n = 4 · exhaustive enumeration of all (EK, QK) pairs

1 Purpose

This document presents an exhaustive empirical run of Enqpy in a degenerate-radix small-n configuration (R = 2, n = 4). The full keyspace of 256 (EK, QK) pairs is enumerated. Each pair is run through Phase 1 through Phase 4 of the cipher against a common public nonce and plaintext, and the resulting ciphertexts are grouped by value.

The purpose is to demonstrate, empirically and non-probabilistically, the {all · valid · one} structure that underlies Enqpy's underdetermination claim:

- **All** — the complete keyspace that an adversary must consider.
- **Valid** — the meaningfully large subset of keys that produce the observed ciphertext.
- **One** — the single correct key, indistinguishable from its equivalents by any analysis of (public nonce, ciphertext, plaintext).

In the MOD16 production cipher this structure is guaranteed by the PDAF Mode 1 preimage bound and the Ideal Configuration construction. At R = 2, the same structural ambiguity can be observed directly by enumeration — no sampling, no statistics, no claims beyond direct computation. Section 10 revisits the same enumeration under the plaintext-equivocation (PET) question from the ePrint paper's Theorem 3: the single 43-partition of the 256 key pairs witnesses both the key-equivocation axis and the plaintext-equivocation axis discussed in the ePrint paper.

2 Setup — Parameters

All quantities in this run are fixed as follows.

Parameter	Value
Radix R	2 (binary — digits are 0 or 1)
Key length n	4 digits per key (EK, QK, VKP, OKP, VKC, OKC all length n)
Tiling length T	24 (= 2n + 16 per FCD §5.2.3; maximum access index is p + δ + 1 + c ≤ 2n + 14)

© 2026 NQP LLC • enqpy.com • Released for evaluation and citation; patents bound open under an irrevocable public covenant

Parameter	Value
Domain separator DS_SEP	0xF (per FCD §4.2 step 3; at R = 2 the analogue is the all-ones digit, DS_SEP mod R = 1)
Plaintext length L	8 digits
Full keyspace	$2^n \times 2^n = 16 \times 16 = 256$ (EK, QK) pairs
W stream length	$3 \cdot n^2 = 48$ digits (three cases per pointer step, CS-ordered)

2.1 Chosen example inputs

```
eff_or  = 1010    (public nonce — appears in clear alongside ciphertext)
PT      = 01100011    (plaintext, 8 binary digits)
EK_true = 1011    (selected as the "true" encryption master key)
QK_true = 0101    (selected as the "true" query master key)
```

The attacker is assumed to know eff_or and PT (known-plaintext setting), and to observe the ciphertext CT. The attacker does not know EK_true or QK_true. The question the enumeration answers is: given eff_or, PT, and CT, how many (EK, QK) pairs are consistent?

3 Algorithm — Ideal Configuration, Phase 1–4

The Ideal Configuration derives session pointers from the public nonce only. Neither EK nor QK appears in VKP or OKP derivation. This is the configuration under which the key-equivocation theorem achieves $H(EK, QK \mid \text{transcript}) \geq \log_2(4) = 2$ bits.

3.1 PDAF Mode 1

```
PDAF_Mode1(VK, OK)[c·n + p] = (VK_t[p + OK_t[p] + 1 + c] + VK_t[p + c])  mod R
where VK_t, OK_t are VK, OK tiled cyclically to length T, for c,p in 0..n−1.
```

3.2 Phase 2 (Ideal Configuration)

```
VKP      = PDAF_Mode1(eff_or, eff_or)[:n]
OKP[i]   = (VKP[i] + DS_SEP)  mod R
VKC      = PDAF_Mode1(EK, VKP)[:n]
OKC      = PDAF_Mode1(QK, OKP)[:n]
```

3.3 Phase 2B — Case Selector

```
CS_raw   = PDAF_Mode1(VKC, OKC)[:3]
CS_val   = (CS_raw[0]·R² + CS_raw[1]·R + CS_raw[2])  mod 6
case_order = CS_PERMUTATION_TABLE[CS_val]
CS_PERMUTATION_TABLE  ⟶ {[1,2,3], [1,3,2], [2,1,3], [2,3,1], [3,1,2], [3,2,1]}
```

© 2026 NQP LLC • enqpy.com • Released for evaluation and citation; patents bound open under an irrevocable public covenant

3.4 Phase 3 — Selection (W stream)

For each (p, c) with p,c ∈ 0..n−1, three candidate values are computed and emitted in CS order, giving 3·n² = 48 output digits:

```
Case 1:  (OKC_t[p] + VKC_t[p+c])  mod R
Case 2:  δ  = OKC_t[p];   (VKC_t[p+δ+1+c]  + VKC_t[p+c]) mod R
Case 3:  δ₂ = VKC_t[p];   (OKC_t[p+δ₂+1+c] + OKC_t[p+c]) mod R
```

3.5 Phase 4 — XOR cipher

```
CT[i]    = (PT[i] + W[i])  mod R      for i = 0..L−1
```

For R = 2, addition mod R is XOR. Phase 5 (in-session key update) is not exercised in this single-message run.

4 Walkthrough — Single Encryption with the True Key

This section traces a single encryption end-to-end using EK_true and QK_true. All intermediate arrays are shown.

4.1 Inputs

```
eff_or  = 1010
EK_true = 1011
QK_true = 0101
PT      = 01100011
```

4.2 Phase 2 — Session key derivation

```
VKP = PDAF_Mode1(eff_or, eff_or)[:n]  = 0101
OKP = (VKP + 1) mod 2                 = 1010
VKC = PDAF_Mode1(EK_true, VKP)[:n]    = 1101
OKC = PDAF_Mode1(QK_true, OKP)[:n]    = 0101
```

4.3 Phase 2B — Case Selector

```
CS_raw      = PDAF_Mode1(VKC, OKC)[:3]  = 001
CS_val      = 1
case_order  = [1, 3, 2]
```

4.4 Phase 3 — W stream

```
W  (length 3·n² = 48 digits, first 8 used):
W[0..47] = 10000001
```

4.5 Phase 4 — Ciphertext

```
CT[i] = PT[i] XOR W[i]   for i = 0..7
```

CT = 11100010

© 2026 NQP LLC • enqpy.com • Released for evaluation and citation; patents bound open under an irrevocable public covenant

The full encryption of plaintext 01100011 under (EK_true, QK_true) = (1011, 0101) with public nonce 1010 yields ciphertext 11100010. The adversary is assumed to observe this ciphertext together with the public nonce; the known-plaintext setting also gives the adversary PT. The question is what (EK, QK) pairs are consistent.

5 Exhaustive Enumeration — Distribution

All $2^4 \cdot 2^4 = 256$ (EK, QK) pairs were run against the fixed (eff_or, PT). Each pair produces a ciphertext. Pairs are grouped by identical ciphertext. The distribution of group sizes is shown below.

Class size	# groups	Pairs covered	Interpretation
4	29	116	Theoretical minimum — $\log_2(4) = 2$ bits residual entropy
8	8	64	Finite-CT slack above theoretical minimum
12	5	60	Finite-CT slack above theoretical minimum
16	1	16	Finite-CT slack above theoretical minimum
TOTAL	**43**	**256**	43 distinct ciphertexts map to 256 (EK,QK) pairs

Key observation: the minimum group size across the full enumeration is 4. No ciphertext in this run is produced by fewer than 4 distinct (EK, QK) pairs. This matches the formal Ideal Configuration theorem, which establishes a lower bound of $\log_2(4) = 2$ bits on H(EK, QK | transcript). The empirical bound is observed directly, without any probabilistic assumption.

Of 256 possible ciphertexts in $\{0,1\}^8$, only 43 occur. The cipher is highly non-injective on (EK, QK) → CT under this particular (eff_or, PT) pair — exactly the property that blocks unique key identification in this toy run for an adversary with any amount of ciphertext.

6 The {all · valid · one} Demonstration

The adversary observes ciphertext CT = 11100010. Under known-plaintext assumptions the adversary also knows the plaintext PT = 01100011 and the public nonce eff_or = 1010. The adversary's task is to determine which (EK, QK) produced the observed ciphertext.

6.1 All — the full keyspace

Every (EK, QK) pair in $\{0,1\}^4 \times \{0,1\}^4$ is a candidate. Size = 256.

© 2026 NQP LLC • enqpy.com • Released for evaluation and citation; patents bound open under an irrevocable public covenant

6.2 Valid — pairs consistent with CT = 11100010

Exhaustive enumeration identifies exactly 4 pairs that produce the observed ciphertext. The adversary cannot distinguish among them using any analysis of (eff_or, PT, CT). The true key is one of these; each of the others is an equally-valid candidate:

#	EK candidate	QK candidate	Status
1	0100	0101	equivalent
2	0100	1010	equivalent
3	**1011**	**0101**	**TRUE KEY**
4	1011	1010	equivalent

6.3 One — the true key

The single correct pair is (EK, QK) = (1011, 0101). It is indistinguishable, by cipher output alone, from the other 3 pairs in the equivalence class.

6.4 Algebraic structure of the class

Under R = 2 the 4 pairs in the class exhibit a clean product structure: they are exactly { EK_a, EK_b } × { QK_a, QK_b }, where each pair of key values is bit-complementary in n = 4. This is the direct algebraic signature of the PDAF Mode 1 many-to-one preimage bound at R = 2 (fiber size = 2), squared across the independent EK and QK derivation paths — yielding the 2 × 2 = 4 minimum ambiguity observed throughout the enumeration.

7 Complete Enumeration — All 43 Ciphertext Groups

Each block below shows one distinct ciphertext and the full list of (EK, QK) pairs that produce it. The true-key group is shaded.

Format: each cell shows EK · QK (both 4-digit binary, space-separated by a dot). An asterisk (*) marks the true key.

Group 1: CT = 00000001 [4 pairs]

0111 · 0001	0111 · 1110	1000 · 0001	1000 · 1110

Group 2: CT = 00000010 [4 pairs]

0001 · 0011	0001 · 1100	1110 · 0011	1110 · 1100

Group 3: CT = 00000011 [4 pairs]

0111 · 0000	0111 · 1111	1000 · 0000	1000 · 1111

Group 4: CT = 00000110 [8 pairs]

© 2026 NQP LLC • enqpy.com • Released for evaluation and citation; patents bound open under an irrevocable public covenant

0010 · 0010	0010 · 0011	0010 · 1100	0010 · 1101
1101 · 0010	1101 · 0011	1101 · 1100	1101 · 1101

Group 5: CT = 00010010 [4 pairs]

0001 · 0010	0001 · 1101	1110 · 0010	1110 · 1101

Group 6: CT = 00010110 [8 pairs]

0101 · 0100	0101 · 0101	0101 · 1010	0101 · 1011
1010 · 0100	1010 · 0101	1010 · 1010	1010 · 1011

Group 7: CT = 00011001 [4 pairs]

0111 · 0101	0111 · 1010	1000 · 0101	1000 · 1010

Group 8: CT = 00011010 [4 pairs]

0011 · 0110	0011 · 1001	1100 · 0110	1100 · 1001

Group 9: CT = 00011011 [4 pairs]

0111 · 0100	0111 · 1011	1000 · 0100	1000 · 1011

Group 10: CT = 00011110 [4 pairs]

0011 · 0111	0011 · 1000	1100 · 0111	1100 · 1000

Group 11: CT = 00100010 [4 pairs]

0100 · 0011	0100 · 1100	1011 · 0011	1011 · 1100

Group 12: CT = 00100100 [4 pairs]

0111 · 0010	0111 · 1101	1000 · 0010	1000 · 1101

Group 13: CT = 00100110 [4 pairs]

0111 · 0011	0111 · 1100	1000 · 0011	1000 · 1100

Group 14: CT = 00110000 [8 pairs]

0100 · 0010	0100 · 1101	0111 · 0110	0111 · 1001
1000 · 0110	1000 · 1001	1011 · 0010	1011 · 1101

Group 15: CT = 00110010 [4 pairs]

0111 · 0111	0111 · 1000	1000 · 0111	1000 · 1000

Group 16: CT = 00111001 [8 pairs]

0010 · 0100	0010 · 0101	0010 · 1010	0010 · 1011
1101 · 0100	1101 · 0101	1101 · 1010	1101 · 1011

Group 17: CT = 01001010 [4 pairs]

0001 · 0000	0001 · 1111	1110 · 0000	1110 · 1111

Group 18: CT = 01010100 [4 pairs]

0100 · 0110	0100 · 1001	1011 · 0110	1011 · 1001

© 2026 NQP LLC • enqpy.com • Released for evaluation and citation; patents bound open under an irrevocable public covenant

Group 19: CT = `01010101` [4 pairs]

0100 · 0000	0100 · 1111	1011 · 0000	1011 · 1111

Group 20: CT = `01010111` [12 pairs]

0000 · 0100	0000 · 0101	0000 · 1010	0000 · 1011
0011 · 0000	0011 · 1111	1100 · 0000	1100 · 1111
1111 · 0100	1111 · 0101	1111 · 1010	1111 · 1011

Group 21: CT = `01011010` [4 pairs]

0001 · 0001	0001 · 1110	1110 · 0001	1110 · 1110

Group 22: CT = `01011100` [4 pairs]

0100 · 0001	0100 · 1110	1011 · 0001	1011 · 1110

Group 23: CT = `01011101` [4 pairs]

0100 · 0111	0100 · 1000	1011 · 0111	1011 · 1000

Group 24: CT = `01011110` [4 pairs]

0011 · 0001	0011 · 1110	1100 · 0001	1100 · 1110

Group 25: CT = `01100011` [12 pairs]

0000 · 0000	0000 · 0001	0000 · 1110	0000 · 1111
0110 · 0011	0110 · 1100	1001 · 0011	1001 · 1100
1111 · 0000	1111 · 0001	1111 · 1110	1111 · 1111

Group 26: CT = `01100111` [4 pairs]

0110 · 0010	0110 · 1101	1001 · 0010	1001 · 1101

Group 27: CT = `01110001` [8 pairs]

0010 · 0000	0010 · 0001	0010 · 1110	0010 · 1111
1101 · 0000	1101 · 0001	1101 · 1110	1101 · 1111

Group 28: CT = `01110011` [12 pairs]

0101 · 0010	0101 · 0011	0101 · 1100	0101 · 1101
0110 · 0110	0110 · 1001	1001 · 0110	1001 · 1001
1010 · 0010	1010 · 0011	1010 · 1100	1010 · 1101

Group 29: CT = `01110111` [12 pairs]

0101 · 0110	0101 · 0111	0101 · 1000	0101 · 1001
0110 · 0111	0110 · 1000	1001 · 0111	1001 · 1000
1010 · 0110	1010 · 0111	1010 · 1000	1010 · 1001

Group 30: CT = `10000110` [4 pairs]

0001 · 0110	0001 · 1001	1110 · 0110	1110 · 1001

© 2026 NQP LLC • enqpy.com • Released for evaluation and citation; patents bound open under an irrevocable public covenant

Group 31: CT = 10001111 [8 pairs]

0010 · 0110	0010 · 0111	0010 · 1000	0010 · 1001
1101 · 0110	1101 · 0111	1101 · 1000	1101 · 1001

Group 32: CT = 10010110 [4 pairs]

0001 · 0111	0001 · 1000	1110 · 0111	1110 · 1000

Group 33: CT = 10100000 [4 pairs]

0011 · 0101	0011 · 1010	1100 · 0101	1100 · 1010

Group 34: CT = 10110010 [4 pairs]

0011 · 0100	0011 · 1011	1100 · 0100	1100 · 1011

Group 35: CT = 11000010 [4 pairs]

0001 · 0101	0001 · 1010	1110 · 0101	1110 · 1010

Group 36: CT = 11000101 [16 pairs]

0000 · 0110	0000 · 0111	0000 · 1000	0000 · 1001
0101 · 0000	0101 · 0001	0101 · 1110	0101 · 1111
1010 · 0000	1010 · 0001	1010 · 1110	1010 · 1111
1111 · 0110	1111 · 0111	1111 · 1000	1111 · 1001

Group 37: CT = 11010010 [4 pairs]

0001 · 0100	0001 · 1011	1110 · 0100	1110 · 1011

Group 38: CT = 11100001 [8 pairs]

0011 · 0011	0011 · 1100	0110 · 0101	0110 · 1010
1001 · 0101	1001 · 1010	1100 · 0011	1100 · 1100

Group 39: CT = 11100010 [4 pairs] ← **TRUE-KEY GROUP**

0100 · 0101	0100 · 1010	**1011 · 0101***	1011 · 1010

Group 40: CT = 11100101 [8 pairs]

0011 · 0010	0011 · 1101	0110 · 0100	0110 · 1011
1001 · 0100	1001 · 1011	1100 · 0010	1100 · 1101

Group 41: CT = 11101011 [4 pairs]

0100 · 0100	0100 · 1011	1011 · 0100	1011 · 1011

Group 42: CT = 11110001 [12 pairs]

0000 · 0010	0000 · 0011	0000 · 1100	0000 · 1101
0110 · 0000	0110 · 1111	1001 · 0000	1001 · 1111
1111 · 0010	1111 · 0011	1111 · 1100	1111 · 1101

Group 43: CT = 11110101 [4 pairs]

© 2026 NQP LLC • enqpy.com • Released for evaluation and citation; patents bound open under an irrevocable public covenant

0110 · 0001	0110 · 1110	1001 · 0001	1001 · 1110

8 Interpretation

Three observations follow from the direct enumeration, independent of any proof framework.

8.1 The lower bound is achieved — not asserted

The Ideal Configuration theorem gives H(EK, QK | transcript) ≥ $\log_2$ (4) = 2 bits. At R = 2, n = 4, this lower bound is not inferred from properties of the proof — it is visible as a floor on class size in the enumeration itself. Every single ciphertext in the 256-pair enumeration is produced by ≥ 4 keys. The empirical floor matches the algebraic bound.

8.2 Scaling to the production cipher

At R = 16 and n = 64 (HIGH profile) the keyspace is 2^{512} rather than 2^8; direct enumeration at production scale is intractable. The asymptotic transcript-level equivocation floor proved by Theorem 2 of the ePrint paper is $\log_2$ (4) = 2 bits at every radix R ≥ 2 — radix-independent, produced by the [+R/2] global shift coset (at R = 2: {EK, ¬EK} × {QK, ¬QK}; at R = 16: {EK, EK+8} × {QK, QK+8}). The toy enumeration exhibits this 2-bit floor directly by observing the 4-element equivalence class across every ciphertext in the 256-pair enumeration. A separate per-session quantity, the Theorem 1 PDAF Mode 1 preimage multiplicity, scales with radix (≥ 2 compatible pairs per OK at R = 2; ≥ 16 at R = 16); that is not the transcript-level residual floor the toy's 4-element class demonstrates. The toy result is not a proof of the production cipher; it is a fully enumerable analogue exhibiting the same claimed pattern of structural underdetermination in a compressed radix regime.

8.3 Why the enumeration is strong supporting evidence

A proof of unsolvability at n = 64 requires trust in the proof. A complete enumeration at n = 4 requires trust only in arithmetic. Taken together, the proof and the toy enumeration provide two independent forms of support for the same claimed structural phenomenon: symbolic at production scale, and arithmetic at toy scale. No single (EK, QK) pair in this run is distinguishable from its equivalents by any amount of analysis of the observable transcript.

9 On the R = 2 Degeneracy

The honest objection to this run is that R = 2 is a degenerate regime of the production cipher. The objection is correct and is addressed directly. This section states precisely what the toy compresses relative to MOD16, what it preserves, and what the enumeration is and is not claiming.

9.1 What the R = 2 regime compresses

© 2026 NQP LLC • enqpy.com • Released for evaluation and citation; patents bound open under an irrevocable public covenant

At the binary radix, three features of the production cipher are narrowed rather than fully exercised:

- **Displacement range.** The PDAF OffsetKey nibble OK_t[p] takes values in {0, 1} at R = 2, versus {0, …, 15} at R = 16. The "dynamic index shift" that produces the entangled constraint structure in PDAF Mode 1 narrows from 16 options per position to 2.
- **Case Selector entropy.** CS_raw is 3 digits. At R = 2 it encodes $2^3 = 8$ possible values, reduced mod 6 to one of 6 permutations. At R = 16 it encodes $16^3 = 4096$ values before the mod-6 reduction — a much denser randomization of the Selection-phase ordering.
- **Per-digit fiber.** The MOD-R fiber is R-to-1: each output digit of a modular combination is consistent with R preimage pairs on its value axis. At R = 2 this is 2; at R = 16 this is 16. Each output digit in the production cipher carries eight times the per-digit structural uncertainty of the toy.

9.2 What the R = 2 regime preserves

The Ideal Configuration theorem establishes $H(EK, QK \mid T_\infty) = \log_2(4) = 2$ bits at every radix $R \geq 2$. The 2-bit residual is structural: it is the joint effect of the EK and QK derivation paths each admitting one axis-wise equivalence, yielding a $2 \times 2 = 4$ minimum equivalence class. This structure is independent of radix.

The R = 2 enumeration observes this floor directly. The minimum class size across all 256 pairs is 4. The structure of the class is precisely {EK, EK'} × {QK, QK'} with each axis-wise equivalent being a bit-flip of the other. The production cipher at R = 16, n = 64 produces the same 4-element asymptotic class — it is only the convergence rate to that class (the "finite-CT slack" above the floor) that differs.

9.3 The MOD16 analogue

Direct enumeration at production scale is intractable: at R = 16, n = 64 the keyspace is $2^{512} \approx 10^{154}$ (EK, QK) pairs, far beyond computational reach. What the toy enumeration cannot do is produce the analogous class table at production scale. What it can do is provide an arithmetic witness to the structural claim the proof makes about that scale, at a scale where the property is checkable by direct enumeration rather than by formal argument.

Property	Toy (R = 2, n = 4)	Production (R = 16, n = 64 HIGH)
Keyspace size	$2^8 = 256$ pairs	2^{512} pairs
PDAF per-digit fiber	2	16
OffsetKey displacement	{0, 1}	{0, …, 15}
CS_raw value space	8 (mod 6 → 6)	4096 (mod 6 → 6)

© 2026 NQP LLC • enqpy.com • Released for evaluation and citation; patents bound open under an irrevocable public covenant

Property	Toy (R = 2, n = 4)	Production (R = 16, n = 64 HIGH)
W stream length	3 · 16 = 48 digits	3 · 4096 = 12,288 digits
Asymptotic class floor	4 pairs (2 bits)	4 pairs (2 bits)
Observed at finite CT (this run)	4 minimum; 4–16 observed	theoretically convergent to 4; not directly enumerable
Method of verification	exhaustive enumeration	formal proof + structural analogue

9.4 What the toy is and is not claiming

The toy is claiming: (a) that the PDAF Mode 1 many-to-one preimage structure is real and algebraic, not probabilistic; (b) that under the Ideal Configuration this structure produces a non-trivial equivalence-class floor on (EK, QK) consistency with observed ciphertext; (c) that the floor predicted by the theorem ($\log_2$ (4) = 2 bits) is observed exactly at R = 2, n = 4 with zero slack below it.

The toy is not claiming: (a) that the R = 2 cipher is secure in any practical sense — the displacement range and case-selector entropy are too compressed for deployment at any radix below MOD16; (b) that enumeration at R = 2 proves security at R = 16 — that remains the domain of the formal proof; (c) that R = 2 exercises the full entanglement structure of the production cipher — it does not, and this is stated plainly here rather than discovered in review.

9.5 How this run supports the formal proof

The formal proof is symbolic. Its conclusion is a structural bound on residual entropy. A reviewer who contests the proof's framing can still be shown the enumeration and asked to locate, by any method, a (EK, QK) pair in this run whose ciphertext is not shared with at least 3 other pairs. No such pair exists. The enumeration and the proof exhibit the same structural pattern at different scales, by different means, and they agree.

Proof and enumeration are independent lines of evidence. Neither relies on the other. The R = 2 toy is not a substitute for the R = 16 proof and is not presented as one — it is a second, arithmetic-only witness to the same claimed structural property. That independence is the reason this section exists.

10 The PET Axis: Same Enumeration, Second Interpretation

Sections 2 through 9 present the enumeration under the key-axis question: which (EK, QK) pairs produce a given ciphertext? This section revisits the same enumeration under the PET-axis question: which plaintexts are consistent with a given ciphertext? The two questions are

© 2026 NQP LLC • enqpy.com • Released for evaluation and citation; patents bound open under an irrevocable public covenant

independent underdetermination axes discussed in the ePrint paper — Theorem 2 establishes the first axis (key equivocation), Theorem 3 establishes the second (plaintext equivocation). A single enumeration at R = 2, n = 4 witnesses both.

10.1 The second question

On the key axis, the adversary observes (eff_or, PT, CT) and asks: which (EK, QK) pair produced this? The answer at R = 2, n = 4 is the 4-element equivalence class enumerated in §6.

On the PET axis, the adversary observes (eff_or, CT) — without the plaintext — and asks: what plaintexts PT' could this ciphertext have encrypted under some key pair (EK', QK')? Formally, S(CT, OR) = {PT' : ∃ (EK', QK') with W(EK', QK', OR) ⊕PT' = CT}. This is the consistent-plaintext set defined in the paper's §12.1. The size of S(CT, OR) is the adversary's residual plaintext ambiguity for the observed ciphertext.

10.2 The algebraic identity

Paper Theorem 3b states S(CT, OR) = CT ⊕Im(K_OR), where Im(K_OR) is the image of the keystream function (EK, QK) → W(EK, QK, OR) under fixed public nonce OR. Immediate consequence: |S(CT, OR)| = |Im(K_OR)| for every CT. The consistent-plaintext count is a property of the keystream image, not of the specific ciphertext observed.

For the toy enumeration this identity is directly verifiable. Every (EK, QK) pair produces one W value under the fixed nonce. The 256 pairs partition by W; the number of distinct W values is |Im(K_OR)|. For each distinct W, the set of pairs producing it is its fiber. The same partition that answers the key-axis question (fiber sizes = equivalence class sizes) answers the PET-axis question (|S| = # of distinct W = # of fibers).

10.3 Enumeration results at R = 2, n = 4

Enumerating (EK, QK) over all 256 pairs under fixed eff_or = 1010: |Im(K_OR)| = 43 distinct W values out of 256 possible 8-bit strings — an image density of 16.80%. The keystream function is highly non-surjective at toy scale: 213 of the 256 possible 8-bit strings are not producible as keystreams under this nonce.

Therefore for the chosen ciphertext CT_true = 11100010, the consistent-plaintext set S(CT_true, eff_or) has exactly 43 elements out of 256 possible plaintexts — $\log_2(43) \approx 5.43$ bits of plaintext equivocation at toy scale. The adversary who observes CT_true and eff_or but not PT_true cannot narrow the plaintext beyond this 43-element set by any analysis of the cipher output.

© 2026 NQP LLC • enqpy.com • Released for evaluation and citation; patents bound open under an irrevocable public covenant

The toy-scale identity predicted by Theorem 3b is verified directly in this enumeration: |S(CT_true, eff_or)| = |Im(K_OR)| = 43. The algebraic identity holds by observation, not assumption.

10.4 Fiber structure and posterior

The fiber-size distribution over the 43 W values is {4 × 29 fibers, 8 × 8 fibers, 12 × 5 fibers, 16 × 1 fiber} — covering 256 (EK, QK) pairs. This is the same partition as the key-axis enumeration in §5: the class-size histogram there is identical because the same 43 cosets underlie both questions.

Under a uniform prior over (EK, QK), the posterior over S(CT, OR) is proportional to fiber size. The fiber sizes take four distinct values {4, 8, 12, 16} — the posterior over the full consistent-plaintext set is non-uniform. This matches the ePrint paper's explicit disclaimer: Theorem 3c proves posterior uniformity over the Case 1 achievable subset only; extension to the full three-case image is stated as Conjecture 1 (§C.3.4). The toy enumeration empirically observes the Conjecture 1 non-uniformity at R = 2, n = 4 — consistent with the paper's stated position, not a finding that contradicts it.

10.5 Scaling to the production cipher

At R = 16, n = 64 (HIGH profile), Theorem 3 establishes $|S(CT, OR)| \geq 2^{128}$ via the minimum image size |Im_VKP|_min = 16^(n/2) derived from the dependency-graph structure of the Ideal Configuration. The toy result of 43 out of 256 at R = 2, n = 4 is a radix-compressed analogue of this image-size floor, not a weakened instance of the production bound.

Under the same framing as §8.2: the production guarantee of $\geq 2^{128}$ consistent plaintexts is a property of the radix-16 dependency-graph structure; the 5.43-bit toy result is a property of the radix-2 structure. Both demonstrate the same pattern — a non-trivial image-size floor on the keystream function, with the consistent-plaintext set inheriting that floor via Theorem 3b — in regimes that differ quantitatively by 123 bits while preserving the same claimed image-floor pattern in a compressed radix regime.

10.6 What the two axes demonstrate together

A single 43-partition of the 256 (EK, QK) pairs supports the observable claim of structural underdetermination on both the key-equivocation axis and the plaintext-equivocation axis discussed in the ePrint paper, at toy scale. On the key axis, the partition gives the 4-element equivalence class that the adversary cannot resolve given (eff_or, PT, CT). On the PET axis, the same partition gives the 43-element consistent-plaintext set that the adversary cannot resolve given (eff_or, CT).

The paper's two-axis argument (§1 of the ePrint paper) is that even if a reviewer attempts to collapse the key-axis equivocation via a quotient-space or equivalent-key argument (§10 of the

© 2026 NQP LLC • enqpy.com • Released for evaluation and citation; patents bound open under an irrevocable public covenant

ePrint paper), the PET axis still holds independently — it does not depend on the [+R/2] coset structure. The toy observation supports this independence: the 43-element consistent-plaintext set arises from the image structure of the keystream function, not from the coset structure of the key space. The two axes are algebraically separate witnesses to the same phenomenon of structural underdetermination.

End of run.

© 2026 NQP LLC • enqpy.com • Released for evaluation and citation; patents bound open under an irrevocable public covenant